The Next Generation Internet Program

Mari Maeda

ITO

ITO

Today's Internet Traffic Makeup

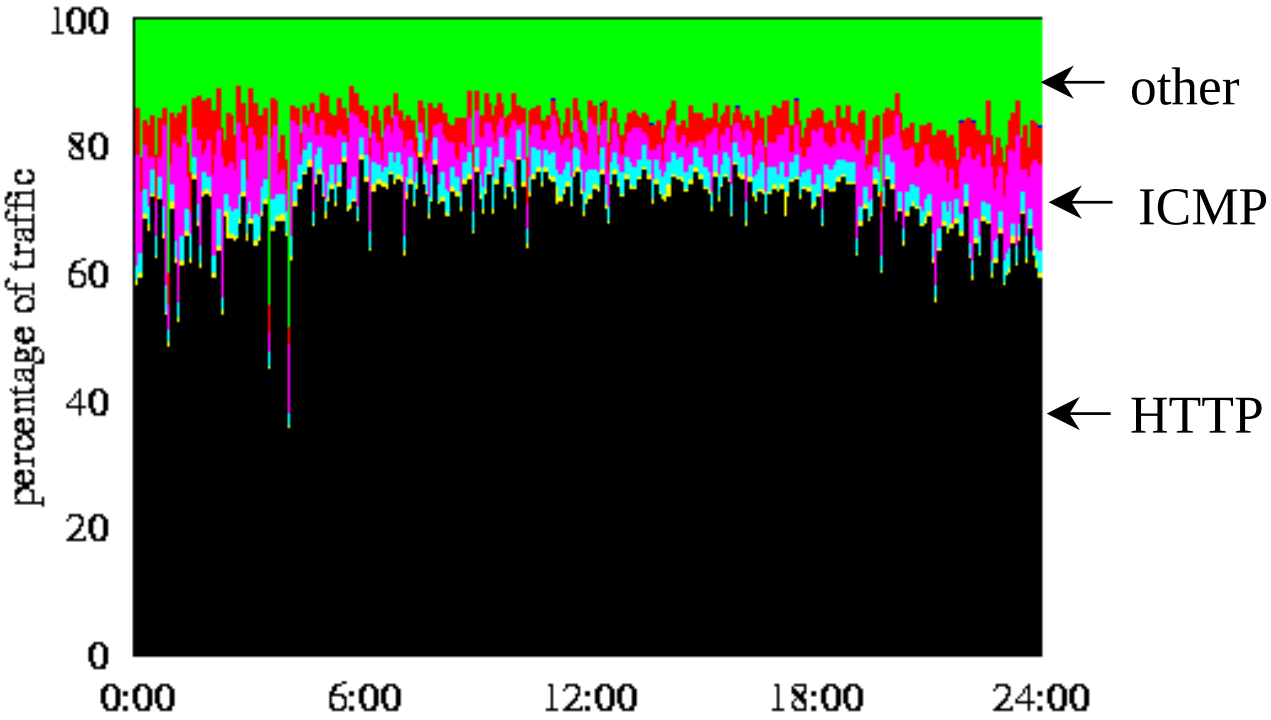

ITO

Today's Internet

Flow Size Distribution

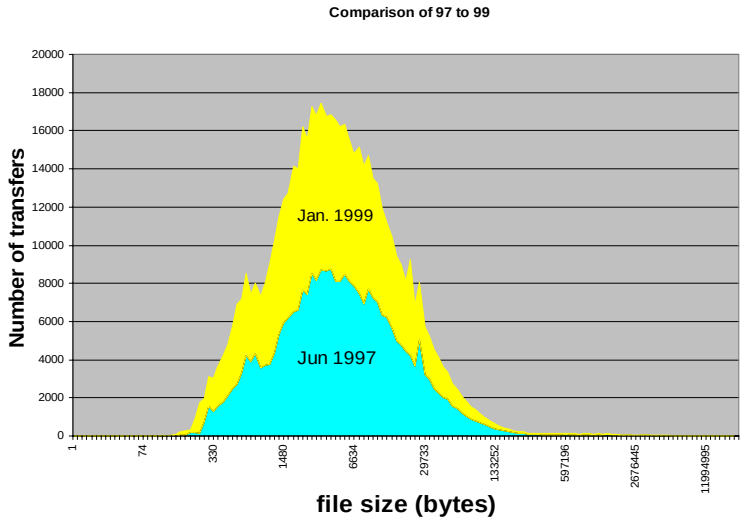

Packet Loss vs. Transmit Rate

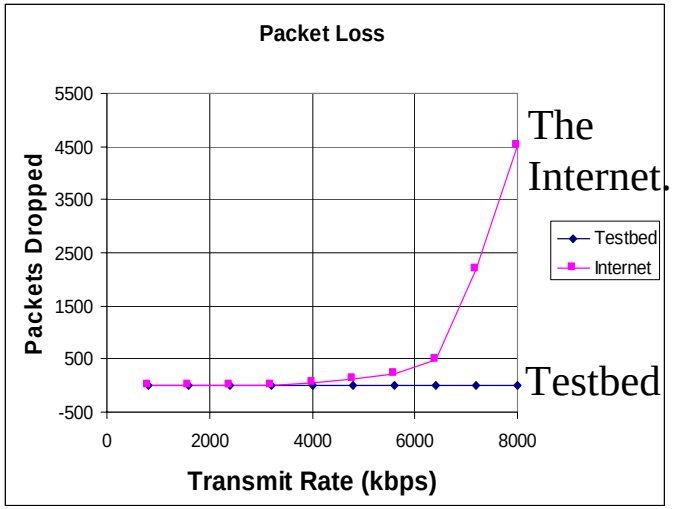

Cambridge to L.A.

Applications

Application binary	10's MB	Digital Video	20-90 Mb
High-Resolution Imagery	100 MB to GB	High-Definition TV	1500 Mbps

ITO

Scaling the Internet

How do we enable the Internet to scale?
(in size,speed,reach,apps)

Number of hosts connected to the Internet

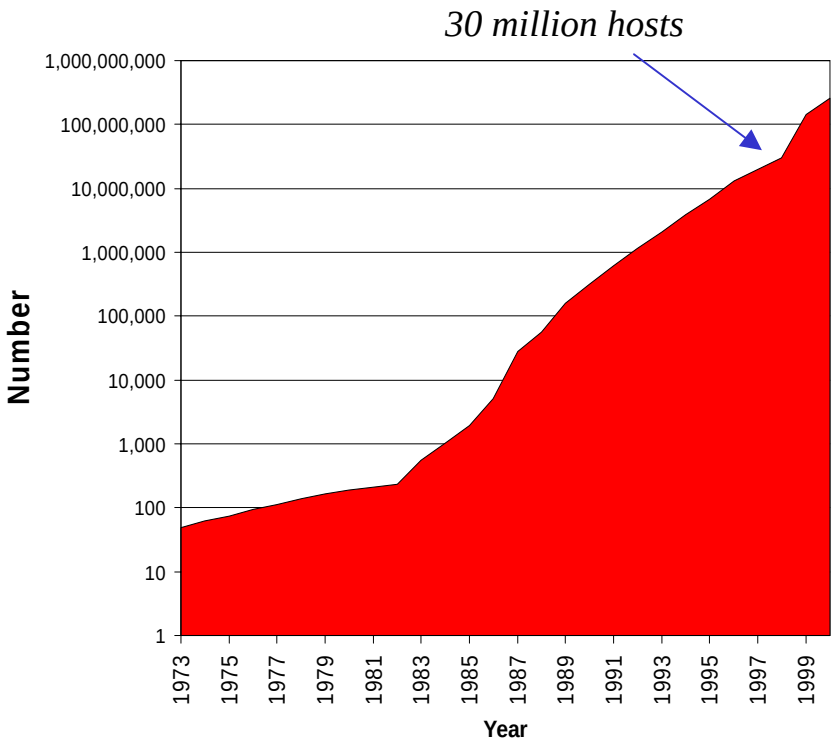

Hop Number Distribution

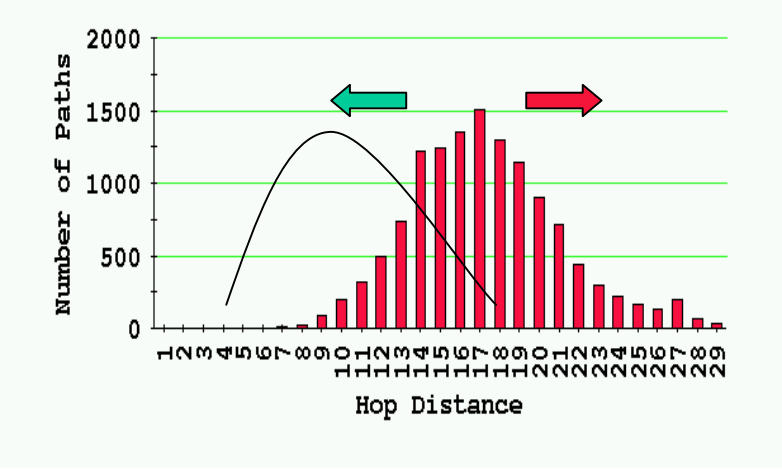

mean hop distance = 16

- Increased loss probab.delay
- delay variation
- decreased security

DARPA's NGI Goals

Develop next generation multiplexing and switching technologies that enable dynamic resource sharing between typical and high-end users

Supernet

Create tools that automate planning and mgmt functions enabling the growth of networks by a factor of 100 or more, while limiting the cost and complexity of network management and control

Network Engineering

ITO

SuperNet Goals

To enable ultra-high bandwidth on demand over national networks, guaranteed over the shared infrastructure

Target: Multi-Gbps end to end

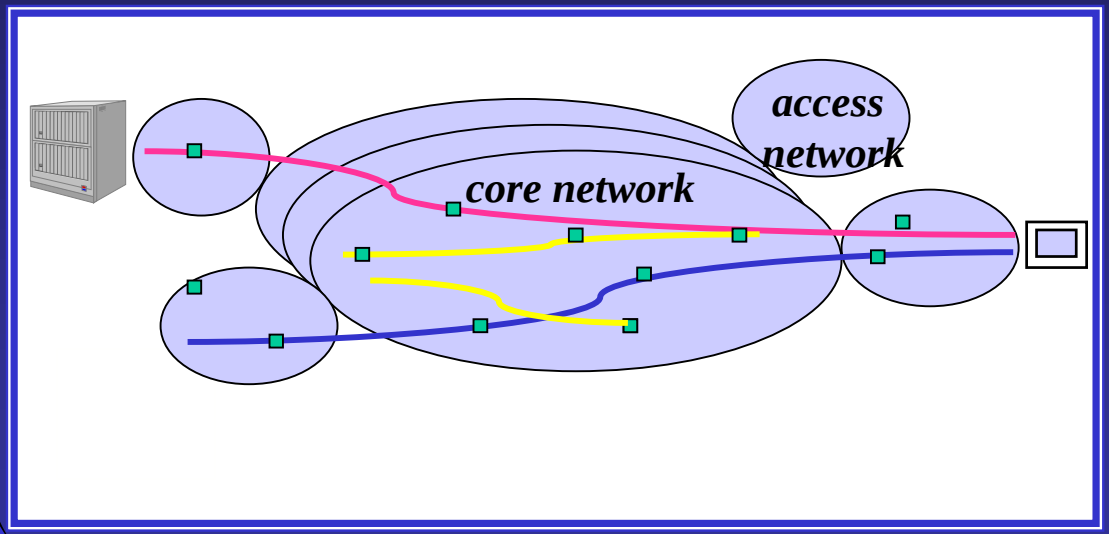

Approach:

- *Streamlined networking protocol stacks*
- *Dynamically reconfigurable/switched optical layer (opaque or electronic)*
- *"Transparency"*
- *New switching/ routing technologies and control algorithms*
- *Dynamic and high bandwidth local access*

ITO

SuperNet: Simplifying Protocol Stacks

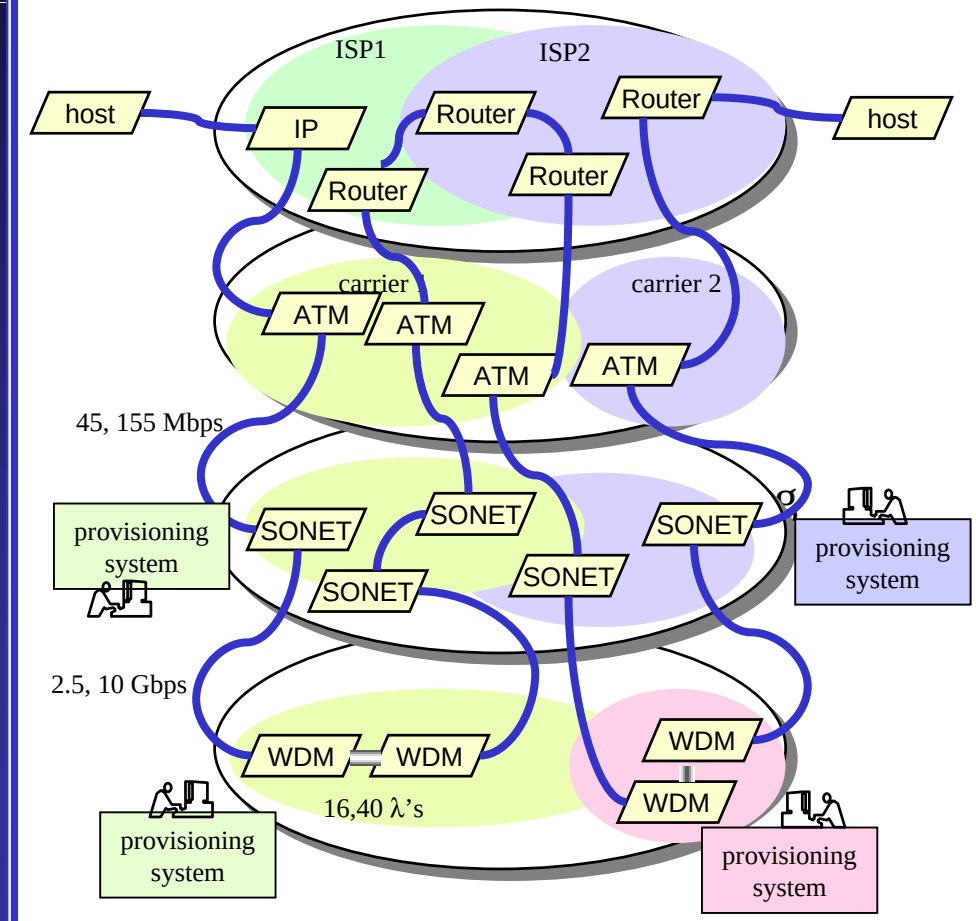

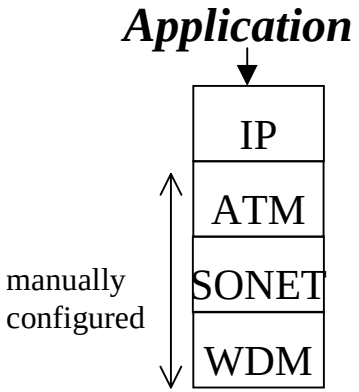

IP over WDM

- WDM based router bypass
- Optical Flow Switching -- based on aggregate traffic change
- Host-triggered path setup
- Optical burst switch (v. short holding times)

speed

Dynamic Optical Layer
transparent, opaque, or
regenerated

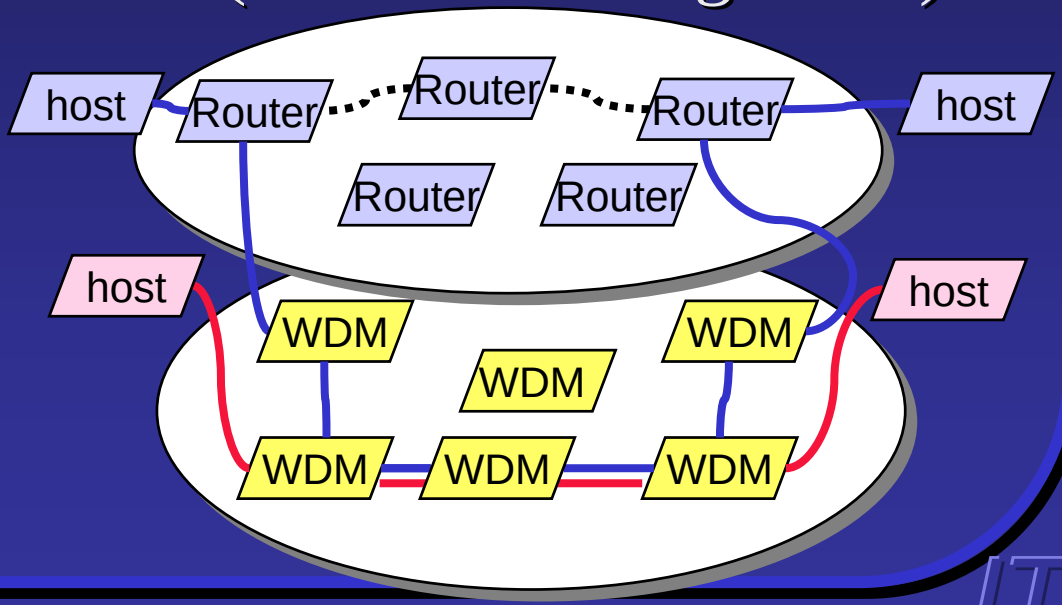

ITO

IP over WDM

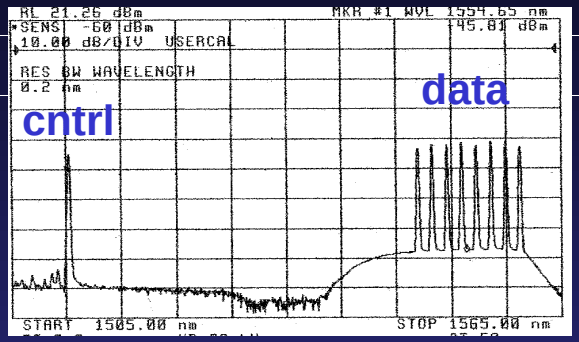

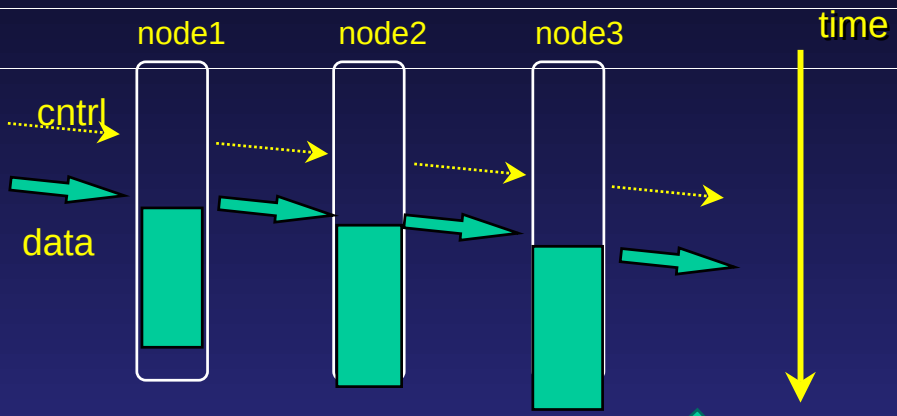

Optical Burst Switch

Optical Label Switching

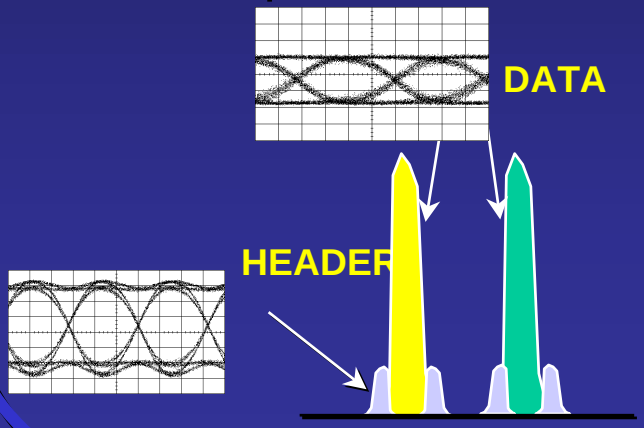

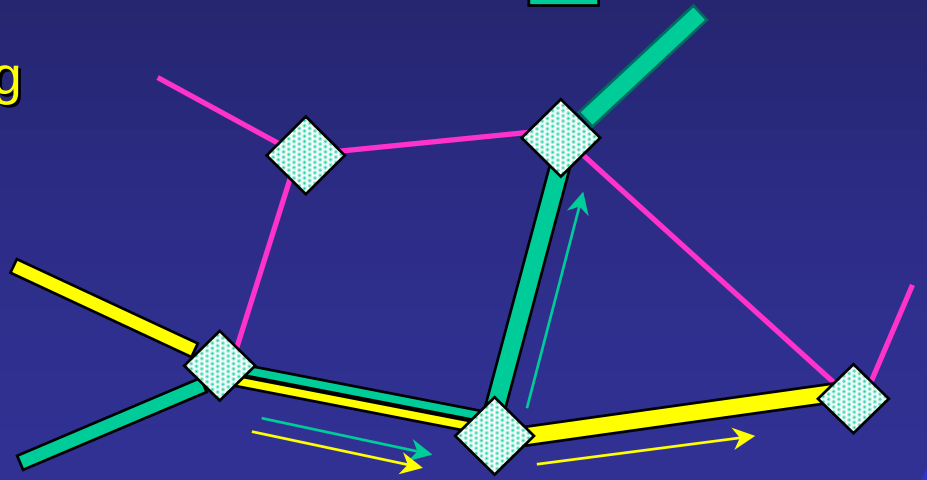

ITO

Bitrate and Protocol Transparent Modules

Modules at the core and the periphery of the network that can

- Recognize and lock to the bit rate (bit-rate adaptability)
- Recognize and handle different protocols (protocol agility)

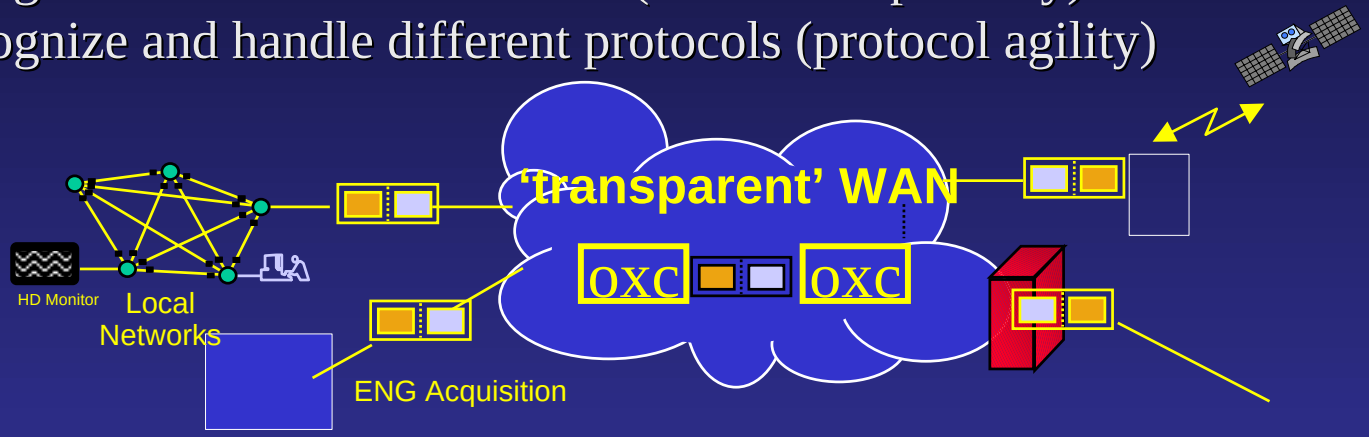

- ***Dynamically reconfigurable or burst switched networks***
- ***Automated network upgrades without replacing hw (lock-on or sw downloads)***
  - ***Rapid deploymet***
  - ***Adapt to new types of sensors, CPE's***
  - ***Minimum inventory***
- ***Development & testing of new protocols***

ITO

Universal Network Access Module

- Target bit range: 100 Mbps to 3 Gbps initially (10 Gbps later)
- Handle a variety of protocol classes at Layer 1 - 3
  - OC3/12/48c ATM / SONET
  - OC3/12/48c IP/SONET
  - Gigabit ethernet
  - SMPTE 25/292
  - IEEE 1394 (firewire)
  - G-Link
  - FDDI
  - Fibre Channel
  - “ngi protocol” e.g. IP/WDM

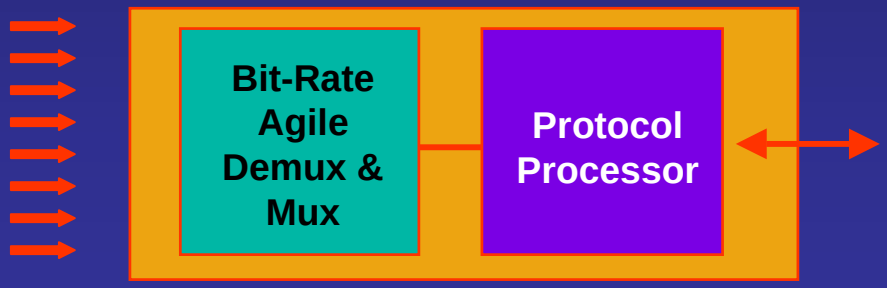

ITO

Network Engineering

- Adaptive control
- Self-management
- Modeling and simulations
- Network visualization

ITO

Network Engineering: Adaptive Network Management Project

Large-scale network fault isolation

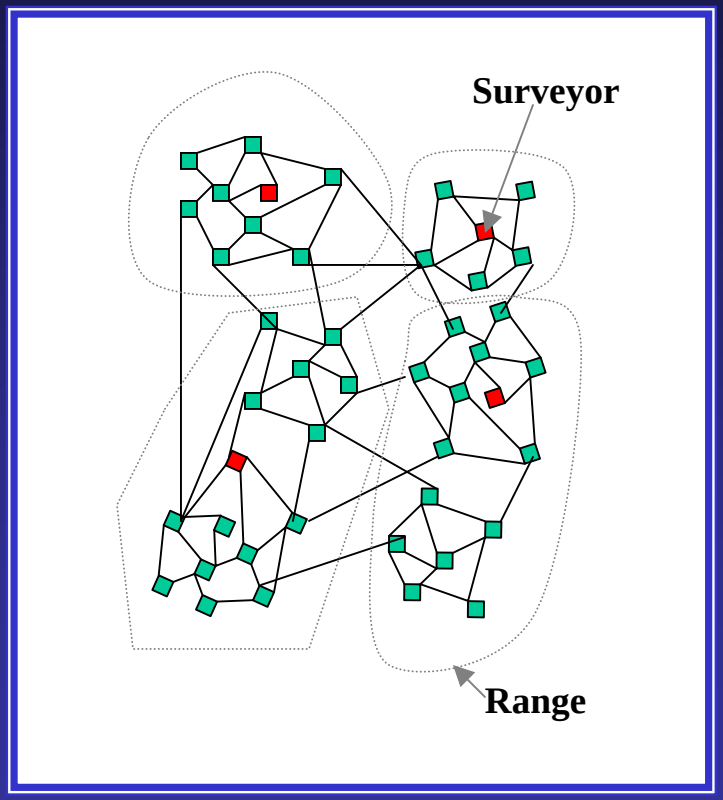

Self-configuring network monitors

- Surveyors map neighborhood
- They coordinate with other surveyors to adjust their ranges
- Careful multicast based self-organization
  - Continuous range expansion
  - Range description exchange
  - Back off
- …eventually adapts to surveyor failure, network partitions

Adapts to network fault (link cut, node failure, congestion, network partition) and surveyor failure.

ITO

Network Engineering: Real-Time Network Simulations

From: Off-line

- Yesterday's traffic situation guides today's provisioning
- Problems fixed after occurrence

To: Realtime

- Live parameter tuning
- Large-scale changes and repair validation prior to fielding

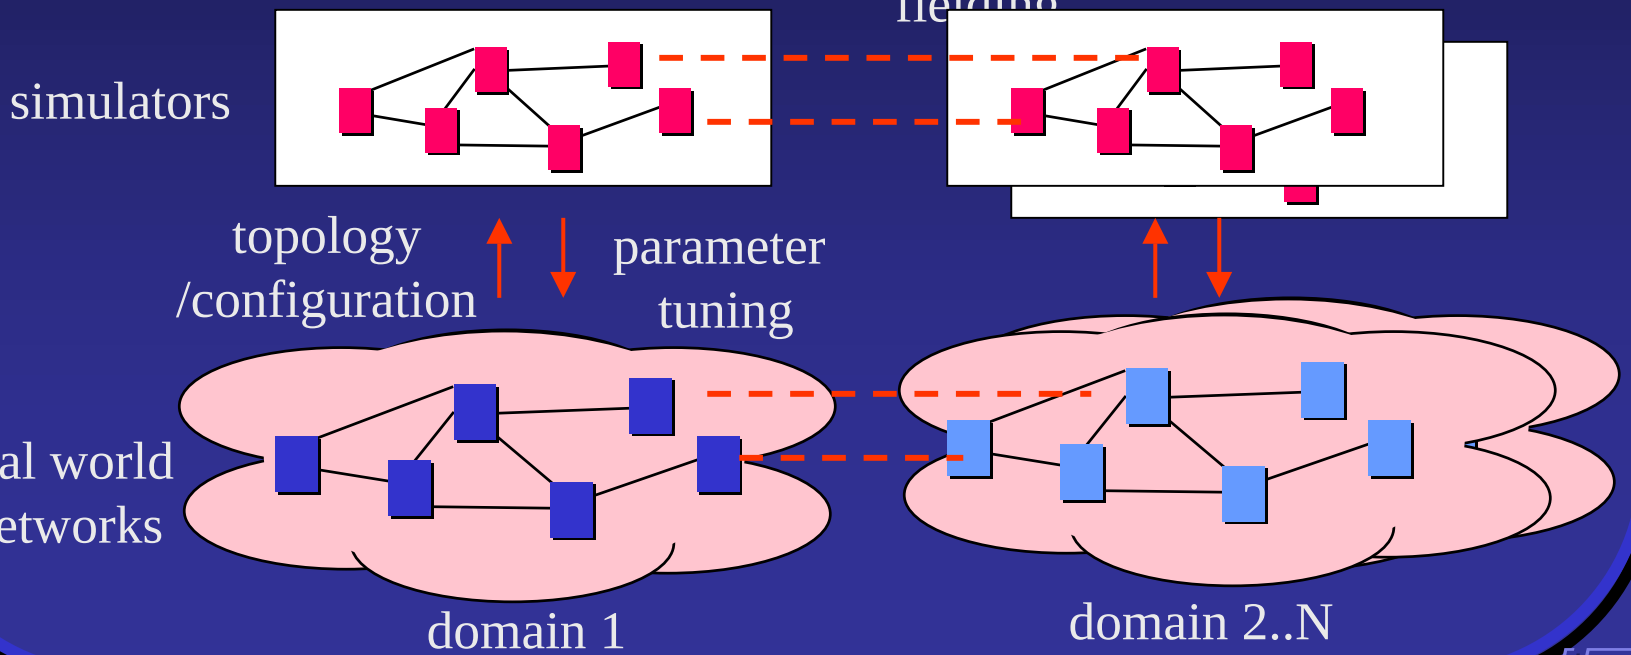

ITO

Adaptive Web Caching Project Target Problem: **"Hot Spots"**

Hundreds of thousands of clients fetching the **same data** from the **same server** at about the **same time**

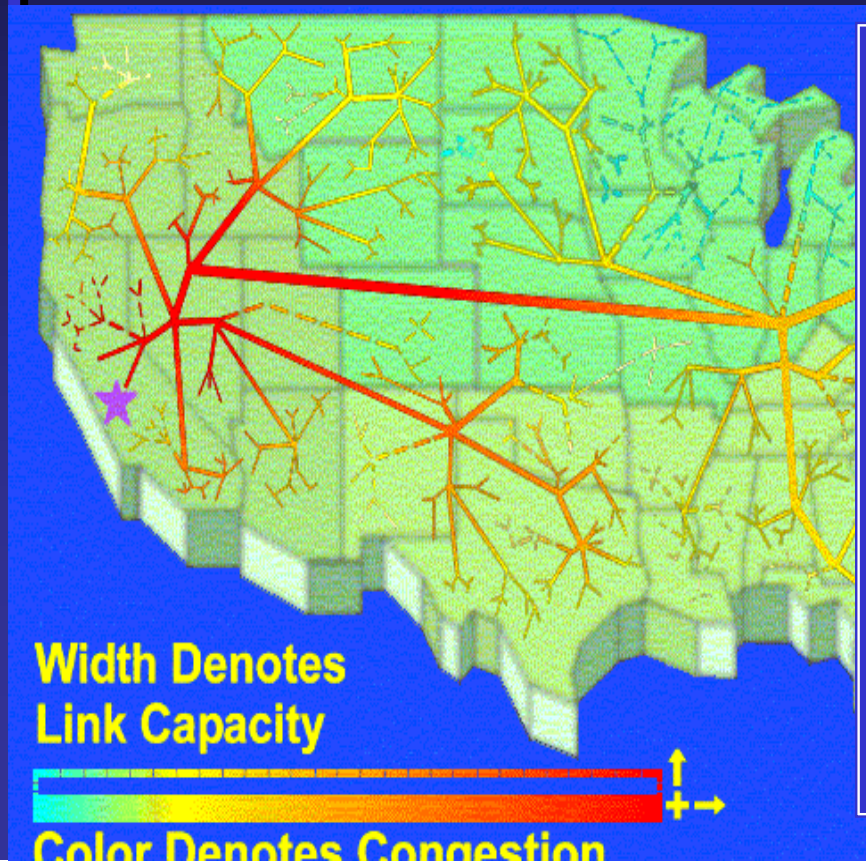

Today:

- Happens few times a year
- Manually create replic. sites
- The Internet has yet to meet the challenge of simultaneous demands from millions of users

Tomorrow:

- Daily occurrence?
- Need demand-driven data dissemination and self-organizing caches e.g. content based routing protocol, cache group management protocol

Network Engineering: Network Monitoring, Analysis and Visualization

- Monitor and automate the discovery of the topology and traffic behavior of the Internet and future networks on a global scale.
- What makes this hard:
  - No central authority
  - Scale (span and speed)
  - Capturing dynamic behavior
  - Visualization

Tools :

"skitter" (active measurements: performance, topology)
"coral" monitors (passive measurements over high speed links)

ITO

Network Tomography

- Network “Radar”: Global connectivity information
- Measure IP paths (“hops”) from source to MANY (~104) destinations
- Use 52 byte ICMP echo requests (every 30 min.) as probes
- Challenges:
  - Pervasive measurement with minimal load on infrastructure
  - Visualization

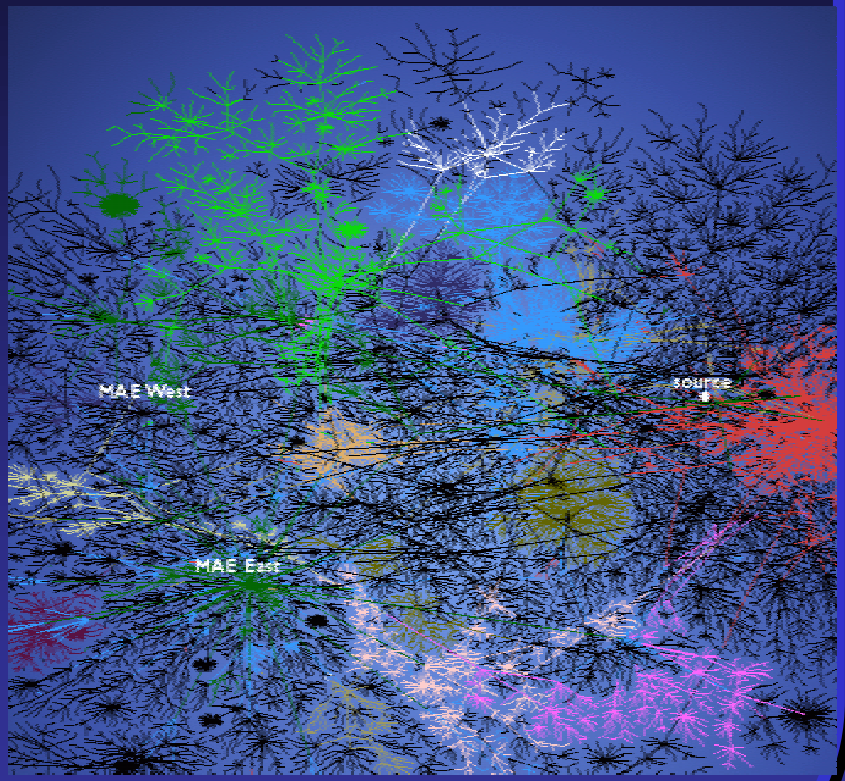

UCSD/CAIDA
(Cooperative Association for Internet Data Analysis)

ITO

Internet Tomography

Hop count histogram

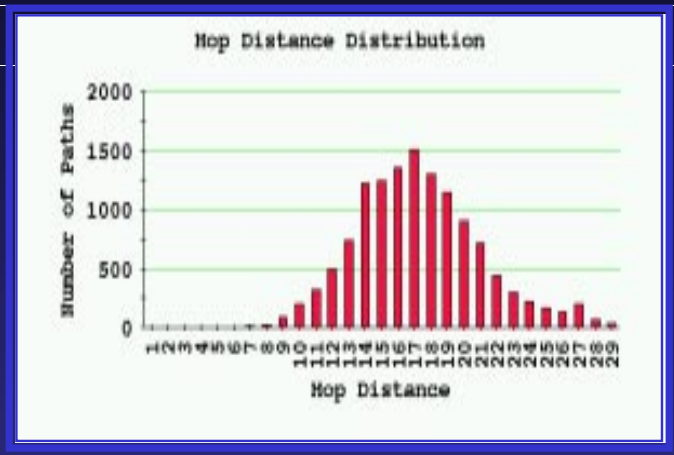

Temporal behavior

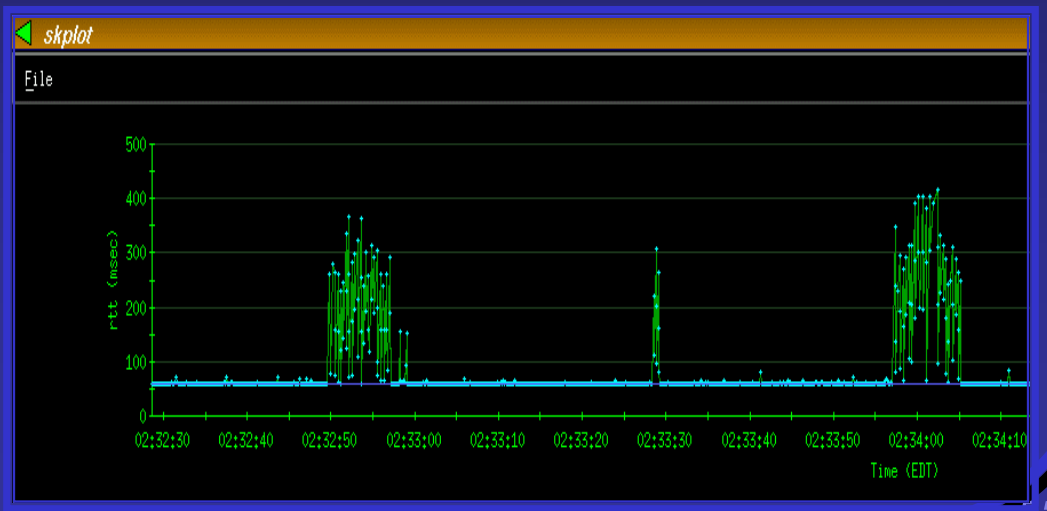

ITO

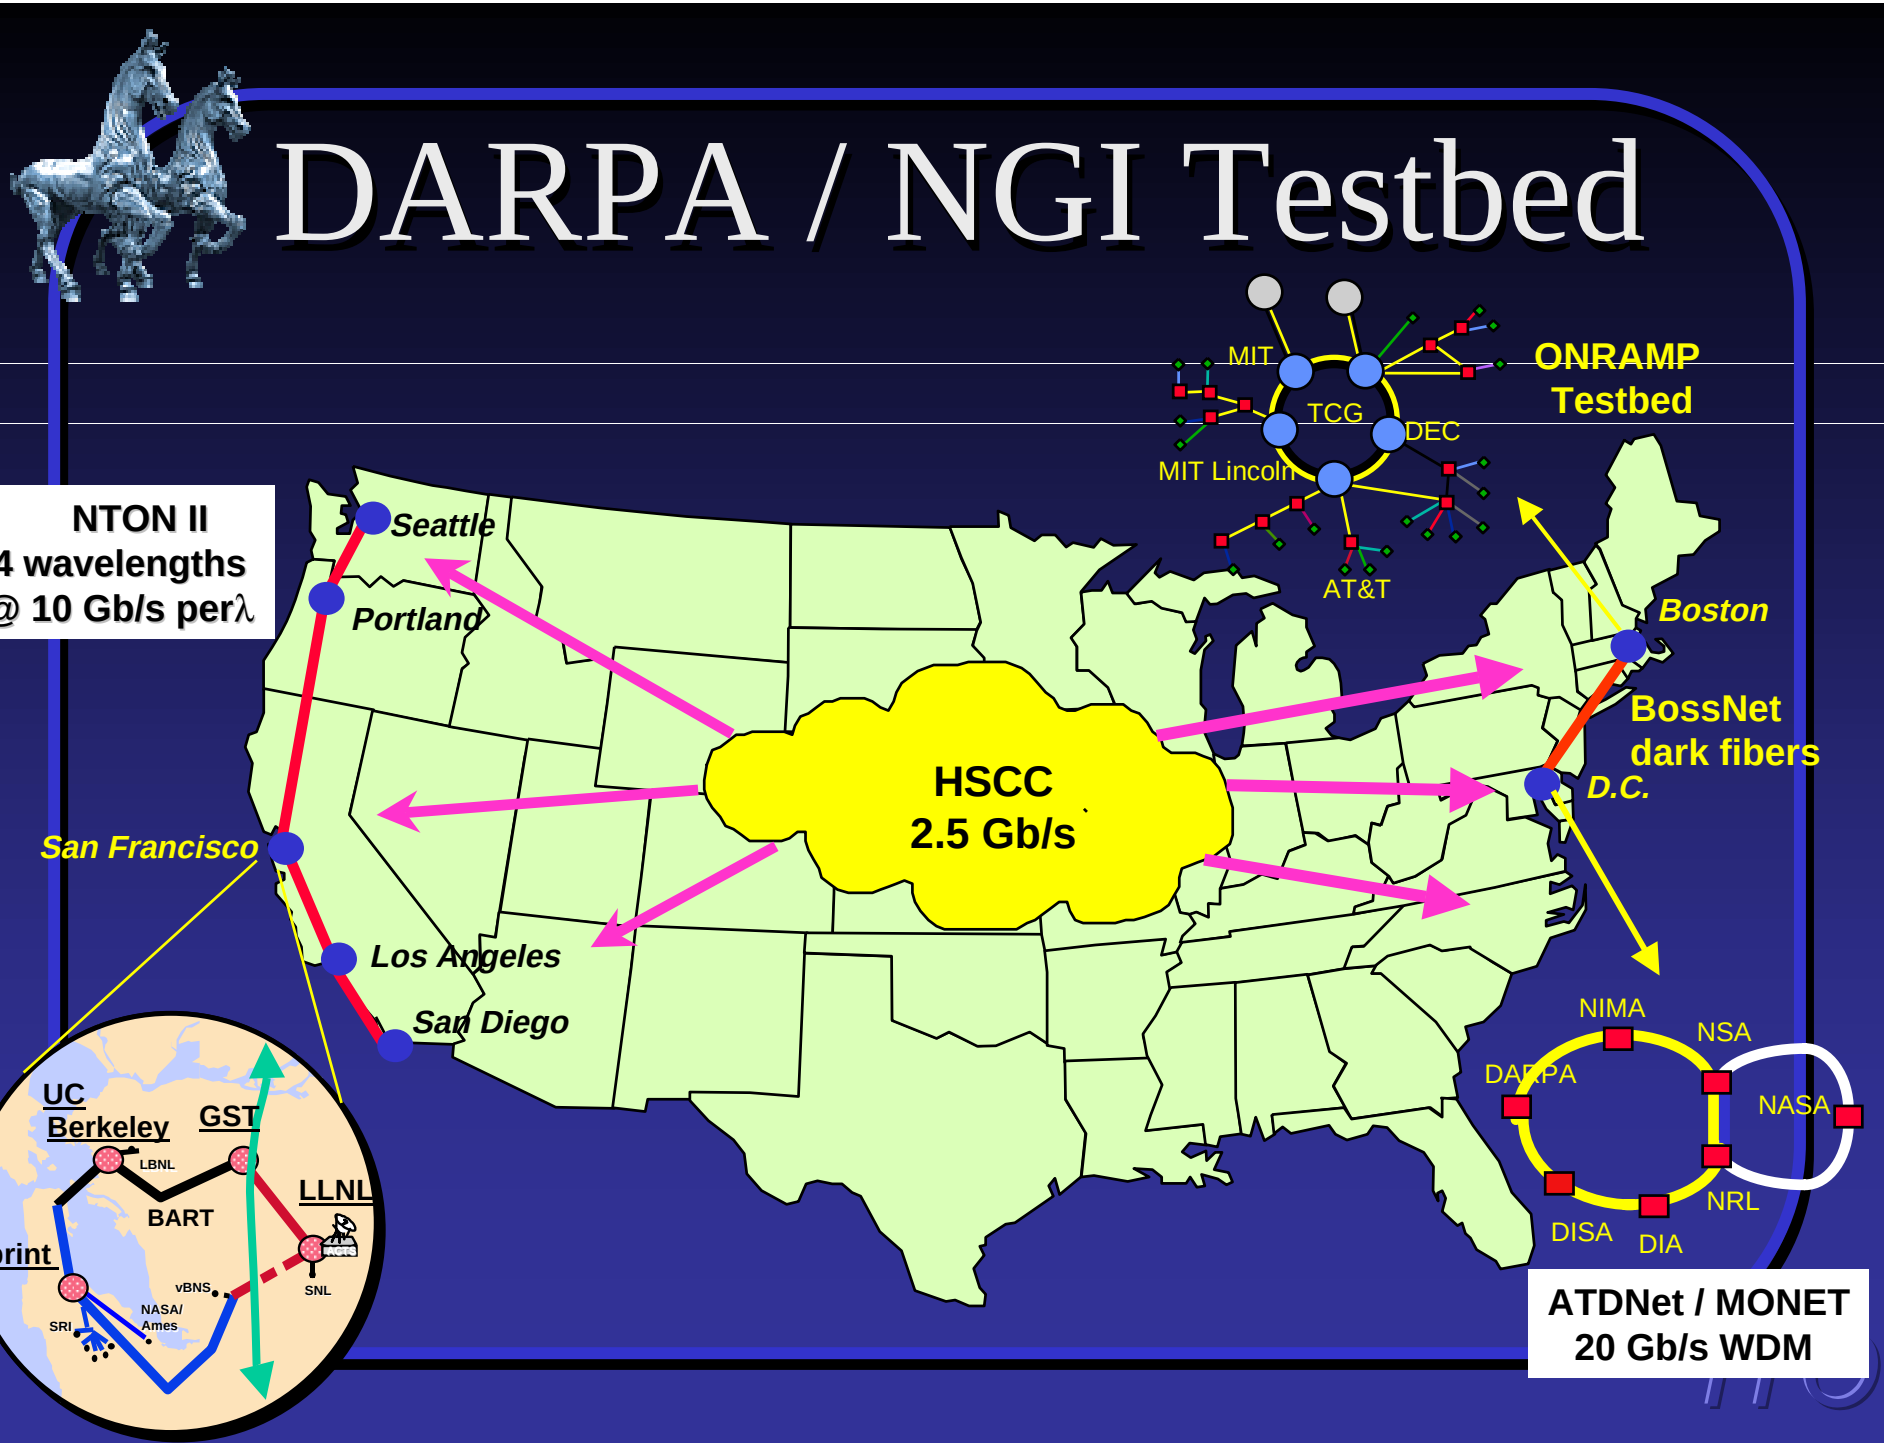
DARPA / NGI Testbed
NTON II
4 wavelengths
@ 10 Gb/s perλ
MIT
TCG
DEC
MIT Lincoln
AT&T
ONRAMP
Testbed
Seattle
Portland
Boston
BossNet
dark fibers
D.C.
HSCC
2.5 Gb/s
San Francisco
Los Angeles
San Diego
NIMA
NSA
DARPA
NASA
NRL
DISA
DIA
UC
Berkeley
GST
LBNL
LLNL
BART
Sprint
vBNS
SNL
SRI
NASA/
Ames
ATDNet / MONET
20 Gb/s WDM

Government-Wide NGI Program

Presidential Initiative -

Start FY1998; 3 year base + 2 year option

Participating Agencies:

DARPA, NSF, NIH/NLM, NIST, NASA, DOE

Goals:

- Networking Research
- Testbeds (SuperNet, vBNS, NREN, ESNET, DREN)
- Revolutionary Applications

ITO